The Expressive Power of Epistemic μ-Calculus

Cătălin Dima[1], Bastien Maubert[2], and Sophie Pinchinat[3]

1 **Université Paris Est, LACL (EA 4219), UPEC, Créteil, France** – dima@u-pec.fr
2 **IRISA, Université de Rennes 1, Rennes, France** – bastien.maubert@irisa.fr
3 **IRISA, Université de Rennes 1, Rennes, France** – sophie.pinchinat@irisa.fr

arXiv:1407.5166v1 [cs.LO] 19 Jul 2014

Abstract

While the μ-calculus notoriously subsumes Alternating-time Temporal Logic (ATL), we show that the *epistemic* μ-calculus does not subsume ATL *with imperfect information* (ATL_i), for the synchronous perfect-recall semantics. To prove this we first establish that jumping parity tree automata (JTA), a recently introduced extension of alternating parity tree automata, are expressively equivalent to the epistemic μ-calculus, and this for any knowledge semantics. Using this result we also show that, for bounded-memory semantics, the epistemic μ-calculus is not more expressive than the standard μ-calculus, and that its satisfiability problem is EXPTIME-complete.

1998 ACM Subject Classification F.4.1 Mathematical Logic, F.4.3 Formal Languages

Keywords and phrases Epistemic μ-calculus, ATL with imperfect information, jumping tree automata, expressiveness

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

The propositional μ-calculus (L_μ) [12] is a logic of utmost importance in theoretical computer science for several main reasons. First, it is a powerful logic that captures all ω-regular properties that are used for the verification of dynamic systems' behavioral properties. In particular, it subsumes all classic temporal logics, such as LTL, CTL and CTL* [6]. Second, it enjoys deep connections with several paradigms that play a fundamental role in modern approaches for the verification of reactive systems: it is equivalent to alternating parity automata [7, Chap. 10], a powerful tool to design decision procedures for temporal logics. L_μ is also closely related with parity games, which are central both for modeling the interaction of systems and for testing the satisfiability of temporal logics [7]. It can be used to specify strategic abilities in multi-player games [16], and it subsumes logics of coalition and strategy like the Alternating-time Temporal Logic (ATL) [1] and Strategy Logic [5]. Finally, its connection with more classic logics is well understood as its expressive power coincides with the bisimulation invariant fragment of the monadic second order logic (MSO) [11].

While most results concern the perfect information setting in which players/agents know the actual state of the system, realistic applications led to consider agents that have to strategize based on a partial information of their environment. This need gave rise to a proliferation of frameworks to represent, reason about and/or strategize under imperfect information. There are basically two trends. One trend relies on extensions of previous strategic logics with additional constraints on strategic abilities of players, that forces them to strategize consistently with their available information. This is the case of variants of ATL with imperfect information like ATL_i, ATL_{ir}, ATLK or ATEL [9, 20, 17] – to cite only a few, see also [3] for a recent survey of the various logics of this type. The other trend is based on extensions of temporal logics with epistemic features, sometimes also combined

© Cătălin Dima, Bastien Maubert and Sophie Pinchinat;
licensed under Creative Commons License CC-BY
Conference title on which this volume is based on.
Editors: Billy Editor and Bill Editors; pp. 1–14

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

with the concepts of the former. Such logics include Epistemic Temporal Logic [8], epistemic mu-calculus L_μ^K, first introduced in [18], and the epistemic alternating mu-calculus AMC [4].

Comparing the two trends is necessary to share expertise, and it is relevant to wonder whether L_μ^K has the same central position as the standard μ-calculus has in the perfect information setting. Some results are already known: the epistemic μ-calculus subsumes Epistemic Temporal Logic and Propositional Epistemic Dynamic Logic [18], and a notion of Alternating Epistemic Mu-Calculus that considers one-step strategic abilities [4]. It is also known from [4] that $\mathrm{ATL_i}$ is not subsumed by this Alternating Epistemic Mu-Calculus for a memoryless semantics with imperfect information.

Our contribution is threefold: first, we show that the epistemic μ-calculus has the same expressive power as the recently introduced *jumping automata*, an extension of alternating parity tree automata that allow for jumps between tree nodes [15]. The proof relies on the classic result that the modal μ-calculus is equivalent with alternating tree automata [7]

Second, combining this general result with the fact that jumping automata equipped with *recognizable* relations between tree nodes translate in linear time into two-way tree automata [15], we obtain two corollaries: for bounded memory semantics, (1) L_μ^K is not more expressive than L_μ, and (2) the satisfiability problem for L_μ^K is Exptime-complete.

Third, we prove that, unlike in the perfect information setting, $\mathrm{ATL_i}$ is not subsumed by the epistemic μ-calculus: we consider the formula $\langle\!\langle a \rangle\!\rangle \mathbf{F} p$, which means that Alice has a uniform strategy (*i.e.* a strategy consistent with her observations) to eventually reach p, and we show that if a jumping automaton accepts all the (tree) models of this formula then it also accepts another model in which Alice only has a non-uniform strategy to achieve $\mathbf{F}p$. This result is proved for the synchronous and perfect recall semantics of indistinguishability.

The paper is organized as follows. In Section 2, first we introduce basic notations and we recall classic parity games as well as game bisimulations. We then expose the epistemic μ-calculus, ATL with imperfect information, and jumping tree automata. In Section 3 we prove that the epistemic μ-calculus is equivalent to jumping tree automata, from which we derive corollaries on the expressivity and the complexity of L_μ^K with bounded memory. Using again the correspondence between L_μ^K and jumping tree automata, we prove in Section 4 that ATL with imperfect information is not expressible in L_μ^K, and we conclude in Section 5, where we also comment on the impact of the results on the relationship between the epistemic μ-calculus and the monadic second order enriched with equal-level predicate (see e.g. [19]).

2 Preliminaries

In this section we set some notations concerning infinite trees and parity games, and we recall the definitions of the three main objects considered in this paper: epistemic μ-calculus, ATL with imperfect information, and jumping tree automata.

A *tree* is a nonempty set $\tau \subseteq \mathbb{N}^*$such that if $x \cdot i \in \tau$, then $x \in \tau$ and $x.j \in \tau$ for all $j < i$, and if $x \in \tau$, there exists $i \in \mathbb{N}$ such that $x \cdot i \in \tau$, and if $x.i$. The elements of τ are called *nodes*, and the empty word ϵ is the *root* of the tree. If $x \cdot i \in \tau$, $x \cdot i$ is a *child* of x. The *arity* of a node x is its number of children, and if every node of some tree t has arity at most k, τ is a *k-ary tree*. Given a node x of a tree τ, we let $Paths_\tau(x)$ (or simply $Paths(x)$) be the set of infinite paths $\pi = x_0 x_1 \dots$ in τ such that $x_0 = x$ and for all i, x_{i+1} is a child of x_i. Also, for a path $\pi = x_0 x_1 \dots$ we let $\pi[i] := x_i$. For two nodes x and y, y is a *descendent* of x (written $x \succeq y$) if x is a prefix of y, or equivalently if y can be found on some path that starts in x. We denote by $\tau{\downarrow_x}$ the subtree of τ rooted in x: $\tau{\downarrow_x} = \{y \mid x \succeq y\}$.

Trees may be labelled with *atomic propositions* from a countably infinite set $\mathcal{AP}$ that

we fix. For a finite subset $AP \subset \mathcal{AP}$ of atomic propositions, an *AP-tree* is a pair $t = (\tau, \ell)$, where τ is a tree and $\ell : \tau \to 2^{AP}$ is a *labelling* of the nodes. A node x in a tree is reached by a finite prefix ρ of a path in $Paths(\epsilon)$, say $\rho_x = x_0 \dots x_n$ with $x_n = x$. We define the *word* of x, written $w(x)$, by $\ell(\epsilon)\ell(x_1)\dots\ell(x_n)$.

For simplicity, we may write $x \in t$ instead of $x \in \tau$. Finally, if $t = (\tau, \ell)$ is an AP-tree, $p \in \mathcal{AP}$ and $S \subseteq \tau$, we define $t[p \to S]$ as the $(AP \cup \{p\})$-tree $t' = (\tau, \ell')$, where $\ell'(x) = \ell(x) \cup \{p\}$ if $x \in S$, and $\ell(x) \setminus \{p\}$ otherwise. In other words, $t[p \to S]$ is the same tree as t, except that we make p hold exactly on nodes in S.

2.1 Parity games and game bisimulation

We define two-player turn-based parity games, that we use to define acceptance of trees by parity tree automata. We also define game bisimulations, recently introduced in [2].

Fix an alphabet Σ. For an infinite word $w = a_0 a_1 \dots \in \Sigma^\omega$ and $i \geq 0$, we let $w[i] := a_i$ and $w[0, i] := a_0 a_1 \dots a_i$. For a finite word $u = a_0 \dots a_{n-1} \in \Sigma^*$, its *length* is $|u| := n$.

We define two-player turn-based parity games: A *parity game arena* is a tuple $G = (V, E, C)$, where V is a set of *positions* partitioned between positions of Eve (V_E) and those of Adam (V_A). Binary relation $E \subseteq V \times V$ is a set of *moves* that we assume total, *i.e.* for all $v \in V$, there is $v' \in V$ such that $(v, v') \in E$. Finally, $C : V \to \mathbb{N}$ is a *colouring function*. A *parity game* $\mathcal{G} = (G, v_0)$ is a game arena $G = (V, E, C)$ together with an initial position $v_0 \in V$. Given a parity game $\mathcal{G} = (G, v_0)$, a *play* $\pi \in V^\omega$ is an infinite sequence of positions such that $\pi[0] = v_0$, and for all $i \geq 0$, $(\pi[i], \pi[i+1]) \in E$. A *partial play* $\rho = v_0 \dots v_n \in V^*$ is a finite prefix of a play and it *ends* in v_n. A *strategy* σ for Eve is a partial function $\sigma : V^* \to V$ such that for all partial play ρ ending in $v \in V_E$, $\sigma(\rho)$ is defined and $(v, \sigma(\rho)) \in E$. A play π *follows* a strategy σ if for all $i \geq 0$ such that $\pi[i] \in V_E$, $\pi[i+1] = \sigma(\pi[0, i])$, and similarly for partial plays. For a parity game $\mathcal{G}$ and a strategy σ for Eve in $\mathcal{G}$, we denote by $\mathrm{Out}(\mathcal{G}, \sigma)$ the set of *outcomes* of σ, that is plays in $\mathcal{G}$ that follow σ. A play π is *winning* for Eve if the least colour seen infinitely often along π is even, otherwise π is winning for Adam. A *winning strategy* for Eve is a strategy whose outcomes are all winning for Eve. Finally, as we only consider winning strategies of Eve, we say that position v of a game arena G is *winning* if Eve has a winning strategy in (G, v).

Berwanger and Kaiser introduce in [2] a notion of bisimulation between parity games and they prove that two bisimilar games are equivalent with regards to the existence of winning strategies [1]. This result will be crucial to establish our nonexpressivity result in Section 4.

▶ **Definition 1.** Let $G = (V, E, C)$ and $G' = (V', E', C')$ be two game arenas. A *bisimulation* between G and G' is a binary relation $Z \subseteq V \times V'$ such that:

Colour Harmony: for all $(v, v') \in Z$, $C(v) = C'(v')$,

Zig: for all $(v, v') \in Z$, if there is $u \in V$ such that $(v, u) \in E$, then there is $u' \in V'$ such that $(v', u') \in E'$ and $(u, u') \in Z$, and

Zag: for all $(v, v') \in Z$, if there is $u' \in V'$ such that $(v', u') \in E'$, then there is $u \in V$ such that $(v, u) \in E$ and $(u, u') \in Z$.

For initial positions $v_0 \in V$ and $v'_0 \in V'$, we say that (G, v_0) is *bisimilar* to (G', v'_0), written $G, v_0 \leftrightarrows G', v'_0$, if there is a bisimulation Z between G and G' such that $(v_0, v'_0) \in Z$.

▶ **Proposition 2** ([2]). *For two game arenas G and G', and two respective positions v and v', if $G, v \leftrightarrows G', v'$, then v is winning in (G, v) if and only if v' is winning in (G', v').*

[1] Note that in [2] the definitions are more general and consider games with imperfect information.

2.2 Epistemic μ-calculus

We fix $Var = \{X, Y, \ldots\}$ a countably infinite set of *second order variables*. Given a finite set of *agents* Ag, the syntax of the epistemic μ-calculus L_μ^K is defined by the following grammar:

$$\varphi ::= X \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \Diamond\varphi \mid K_i\varphi \mid \mu X.\varphi(X)$$

where $X \in Var$, $p \in \mathcal{AP}$, $i \in Ag$, and in the last rule X appears only positively (under an even number of negations) in $\varphi(X)$. For a *finite* set of atomic propositions $AP \subset \mathcal{AP}$, we denote by $\mathsf{L}_\mu^K(AP, Ag)$, or simply L_μ^K when the parameters are irrelevant, the set of formulas of the epistemic μ-calculus that only use atomic propositions in AP and agents in Ag.

A model of a formula in $\mathsf{L}_\mu^K(AP, Ag)$ consists in an AP-tree t together with a set of binary relations $\{\rightsquigarrow_i\}_{i\in Ag}$ over $(2^{AP})^*$. In the following, for two nodes x and y in t, $x \rightsquigarrow_i y$ stands for $w(x) \rightsquigarrow_i w(y)$: two nodes are related by $\rightsquigarrow_i$ if their node words are related by $\rightsquigarrow_i$. Intuitively, $x \rightsquigarrow_i y$ means that when the current node is x, Agent i considers possible (up to her knowledge) that node y is the current node. Notice that the relation $\rightsquigarrow_i$ is arbitrary and not necessarily an equivalence relation, as often assumed in epistemic logic. From now on, whenever Ag is clear from the context, $\{\rightsquigarrow\}$ will denote a *relation profile* $\{\rightsquigarrow_i\}_{i\in Ag}$. Finally, interpreting a formula requires a *valuation* $V : Var \to 2^t$; also, given $X \in Var$ and $S \subseteq t$, $V[S/X]$ is the valuation that maps X to S, and is equal to V on all other variables.

The semantics of a formula $\varphi \in \mathsf{L}_\mu^K(AP, Ag)$ on an AP-tree $t = (\tau, \ell)$ with relation profile $\{\rightsquigarrow\}$ and valuation V is the set of nodes $[\![\varphi]\!]_V^t \subseteq t$ defined as follows:

- $[\![X]\!]_V^t = V(X)$
- $[\![p]\!]_V^t = \{x \in t \mid p \in \ell(x)\}$
- $[\![\neg\varphi]\!]_V^t = t \setminus [\![\varphi]\!]_V^t$
- $[\![\varphi \vee \psi]\!]_V^t = [\![\varphi]\!]_V^t \cup [\![\psi]\!]_V^t$
- $[\![\Diamond\varphi]\!]_V^t = \{x \in t \mid x \cdot i \in [\![\varphi]\!]_V^t \text{ for some } i \in [k]\}$
- $[\![K_i\varphi]\!]_V^t = \{x \in t \mid y \in [\![\varphi]\!]_V^t \text{ for all } y \text{ such that } x \rightsquigarrow_i y\}$
- $[\![\mu X.\varphi(X)]\!]_V^t = \bigcap\{S \subseteq t \mid [\![\varphi(X)]\!]_{V[S/X]}^t \subseteq S\}$

Classically, for each formula $\mu X.\varphi(X)$ in L_μ^K, the fact that X appears only positively in $\varphi(X)$ ensures that $S \mapsto [\![\varphi(X)]\!]_{V[S/X]}^t$ is a monotone function, and hence that its least fixpoint exists. $[\![\mu X.\varphi(X)]\!]_V^t$ is defined to be this fixpoint.

If $\varphi \in \mathsf{L}_\mu^K$ is a sentence, *i.e.* it has no free variables, its semantics is independent on the valuation, that we may omit from the semantics. For a sentence $\varphi \in \mathsf{L}_\mu^K$, a relation profile $\{\rightsquigarrow\}$ and a tree t, we write $t, \{\rightsquigarrow\} \models \varphi$ for $\epsilon \in [\![\varphi]\!]_{\{\rightsquigarrow\}}^t$, and we let $\mathcal{L}(\varphi, \{\rightsquigarrow\}) := \{t \mid t, \{\rightsquigarrow\}, \epsilon \models \varphi\}$. Finally, we let L_μ denote the sublanguage of L_μ^K obtained by removing the modalities K_i, and simply write $t, \epsilon \models \varphi$ as relation profile do not play any role in the semantics of L_μ-formulas; thus, for $\varphi \in \mathsf{L}_\mu$ we may use $\mathcal{L}(\varphi) = \{t \mid t, \epsilon \models \varphi\}$.

2.3 Alternating-time Temporal Logic with imperfect information

We now recall the syntax and semantics of Alternating-time Temporal Logic with imperfect information ($\mathrm{ATL_i}$). Again, let Ag be a nonempty finite set of agents. The syntax of $\mathrm{ATL_i}(Ag)$ is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle\!\langle A \rangle\!\rangle \mathbf{X}\varphi \mid \langle\!\langle A \rangle\!\rangle \varphi \mathbf{U} \varphi$$

where $p \in \mathcal{AP}$ and $A \subseteq Ag$.

The semantics of $\mathrm{ATL_i}$ is usually defined on concurrent game structures (see [1]). These are transition systems with states labelled by valuations over some finite set of propositions

AP, and where every transition is labelled by a compound action $a = (a_1, \ldots, a_k)$, which is interpreted as Agent $i \in Ag$ playing action a_i during this transition. The imperfect information is usually introduced by letting each agent observe only a subset of AP, and by deciding whether agents remember the past during a play or not. This induces, for each agent, an equivalence relation between finite plays.

In order to make the comparison with epistemic μ-calculus easier, we instead define the semantics of $\mathrm{ATL_i}$ on what we call tree-arenas:

▶ **Definition 3.** Let $AP \subset \mathcal{AP}$ be a *finite* set of atomic propositions, and for each $i \in Ag$, let Act_i be a nonempty finite set of actions available to Agent i. Define $\mathrm{Act} := \bigtimes_{i \in Ag} \mathrm{Act}_i$, and let $AP_{\mathrm{act}} := \{p_a \mid a \in \mathrm{Act}\}$ where each p_a is an atomic proposition not in AP. An (AP, Act)*-tree-arena* is an $(AP \cup AP_{\mathrm{act}})$-tree $t = (\tau, \ell)$ such that $\ell(\epsilon) \cap AP_{\mathrm{act}} = \emptyset$, and for all $x \in \tau \setminus \{\epsilon\}$, $\ell(x) \cap AP_{\mathrm{act}}$ is a singleton.

For the rest of this section, we fix a finite set $AP \subset \mathcal{AP}$ and a finite set of actions Act_i for each agent $i \in Ag$. For an (AP, Act)-tree-arena $t = (\tau, \ell)$ and a node $x \in \tau$, we write $\ell(x) = (v, a)$, where $a \in \mathrm{Act}$ is the unique (compound) action such that $p_a \in \ell(x)$, and $v = \ell(x) \setminus \{p_a\}$. In addition, given $a = (a_1, \ldots, a_k) \in \mathrm{Act}$, a^i will denote a_i. Note that a tree-arena t can be seen as a concurrent game structure: take a node $x \in t$, and let (v, a) be its label. Node x can be seen as a state of a transition system, v as its label, and a as the label of the only transition reaching x. Concerning the imperfect information, similarly to the previous section, we introduce agents' uncertainty by means of binary relations $\{\sim_i\}_{i \in Ag}$ over $(2^{AP \cup AP_{\mathrm{act}}})^*$. Conversely, the unfolding of every concurrent game structure with imperfect information can be seen as a tree-arena equipped with a relation profile. We now adapt the classic semantics of ATL to our setting.

First we need a few more definitions. Fix an (AP, Act)-tree-arena t and a relation profile $\{\sim\}$. A *strategy* for Agent i is a function $\sigma_i : t \to \mathrm{Act}_i$, that defines the strategic choice of Agent i in each possible situation. Because agents have imperfect information, we classically require strategies to be consistent with the information of the agent: if σ_i is a strategy for Agent i, we require that for each $x, y \in t$ such that $x \sim_i y$, $\sigma_i(x) = \sigma_i(y)$ (note that strategies satisfying this requirement are sometimes called *uniform* strategies [10]). For $A \subseteq Ag$, we call A*-profile* a tuple $\sigma_A = (\sigma_i)_{i \in A}$ where σ_i is a strategy for Agent i, and given an A-profile σ_A and $i \in A$, we let σ_A^i denote the strategy of agent i in σ_A. The *outcome* of an A-profile σ_A for some $A \subseteq Ag$ is the set of behaviours that follow the strategies in the profile, defined as follows. For a node x of t, $\mathrm{Out}(x, \sigma_A) \subseteq Paths(x)$ is the set of paths π in t that start in x and such that for all $k \geq 0$, if (v, a) is the label of $\pi[k+1]$, then $\sigma_A^i(\pi[k]) = a^i$ for all $i \in A$.

The semantics of an $\mathrm{ATL_i}$-formula φ with atomic propositions in AP is given with respect to an (AP, Act)-tree-arena $t = (\tau, \ell)$, a relation profile $\{\sim\}$ and a node $x \in t$:

- $t, \{\sim\}, x \models p$ if $p \in v$, where $(v, a) = \ell(x)$
- $t, \{\sim\}, x \models \neg\varphi$ if $t, \{\sim\}, x \not\models \varphi$
- $t, \{\sim\}, x \models \varphi \vee \psi$ if $t, \{\sim\}, x \models \varphi$ or $t, \{\sim\}, x \models \psi$
- $t, \{\sim\}, x \models \langle\!\langle A \rangle\!\rangle \mathbf{X}\varphi$ if there is an A-profile σ_A such that:
  for all $y \in t$, if $x \sim_i y$ for some $i \in A$, then for all $\pi \in \mathrm{Out}(y, \sigma_A)$,
  $t, \{\sim\}, \pi[1] \models \varphi$
- $t, \{\sim\}, x \models \langle\!\langle A \rangle\!\rangle \varphi \mathbf{U} \psi$ if there is an A-profile σ_A such that:
  for all $y \in t$, if $x \sim_i y$ for some $i \in A$, then for all $\pi \in \mathrm{Out}(y, \sigma_A)$,
  there is $i \geq 0$ such that $t, \{\sim\}, \pi[i] \models \psi$, and for all $0 \leq j < i$, $t, \{\sim\}, \pi[j] \models \varphi$

We define the following classic shorthands: $\top := p \vee \neg p$, and $\langle\langle A\rangle\rangle \mathbf{F}\varphi := \langle\langle A\rangle\rangle \top \mathbf{U} \varphi$. Finally, for a formula $\varphi \in \mathrm{ATL_i}$, a set of (compound) actions Act and a relation profile $\{\rightsquigarrow\}$, we let $\mathcal{L}(\varphi, \mathrm{Act}, \{\rightsquigarrow\}) := \{t \mid t \text{ is a } (\mathit{Free}(\varphi), \mathrm{Act})\text{-tree-arena s.t. } t, \{\rightsquigarrow\}, \epsilon \models \varphi\}$.

▶ Remark. We consider here the most restrictive notion of "having a strategy", *i.e.* having a strategy *"de re"* [10]. However, the result that we prove in Section 4 still holds with less restrictive notions of strategies: *"de dicto"* strategies, or simply uniform strategies

2.4 Jumping tree automata

Jumping tree automata (JTA) were introduced in [15, 14]. Let Ag be a finite set of agents. For a set X, $\mathbb{B}^+(X)$ is the set of positive boolean formulas over X, *i.e.* formulas built with elements of X as atomic propositions and using only connectives $\vee$ and $\wedge$. We also allow for formulas $\top$ and $\bot$, and $\wedge$ has precedence over $\vee$. Elements of $\mathbb{B}^+(X)$ are denoted by $\alpha, \beta \ldots$

▶ **Definition 4.** Let $\mathit{Dir} = \{\Diamond, \Box\} \cup \bigcup_{i \in Ag} \{\diamondsuit_i, \boxminus_i\}$ be the set of *automaton directions*. A *jumping automaton* is a tuple $\mathcal{A} = (AP, Q, \delta, q_0, C)$ where $AP \subset \mathcal{AP}$ is a finite set of atomic propositions, Q a finite set of states, $q_0 \in Q$ an initial state, $C : Q \to \mathbb{N}$ a colouring function, and $\delta : Q \times 2^{AP} \to \mathbb{B}^+(\mathit{Dir} \times Q)$ a transition function.

Let $\mathcal{A}$ be a JTA over AP. The meaning of the *jump directions* $\diamondsuit_i, \boxminus_i$ is given by a relation profile $\{\rightsquigarrow\} = \{\rightsquigarrow_i\}_{i \in Ag}$, where for each i, $\rightsquigarrow_i \subseteq (2^{AP})^* \times (2^{AP})^*$. The acceptance of an input tree $t = (\tau, \ell)$ by $\mathcal{A}$ equipped with a relation profile $\{\rightsquigarrow\}$ is defined on a two-player parity game between Eve (the proponent) and Adam (the opponent): let $t = (\tau, \ell)$ be an AP-tree, and let $\mathcal{A} = (\Sigma, Q, \delta, q_0, C)$. We define the game $\mathcal{G}^{\mathcal{A}}_{t,\{\rightsquigarrow\}} = (V, E, C', v_0)$: the set of positions is $V = \tau \times Q \times \mathbb{B}^+(\mathit{Dir} \times Q)$, the initial position is $(\epsilon, q_0, \delta(q_0, \ell(\epsilon)))$, and a position (x, q, α) belongs to Eve if α is of the form $\alpha_1 \vee \alpha_2$, $[\Diamond, q']$ or $[\diamondsuit_i, q']$; otherwise it belongs to Adam. The possible moves in $\mathcal{G}^{\mathcal{A}}_{t,\{\rightsquigarrow\}}$ are the following:

$$(x, q, \alpha_1 \dagger \alpha_2) \to (x, q, \alpha_i) \qquad \text{where } \dagger \in \{\vee, \wedge\} \text{ and } i \in \{1, 2\} \tag{1}$$

$$(x, q, [\bigcirc, q']) \to (y, q', \delta(q', \ell(y))) \quad \text{where } \bigcirc \in \{\Diamond, \Box\} \text{ and } y \text{ is a child of } x \tag{2}$$

$$(x, q, [\ominus_i, q']) \to (y, q', \delta(q', \ell(y))) \qquad \text{where } \ominus_i \in \{\diamondsuit_i, \boxminus_i\} \text{ and } x \rightsquigarrow_i y \tag{3}$$

Positions of the form $(x, q, \top)$ and $(x, q, \bot)$ are deadlocks, winning for Eve and Adam respectively. The colouring function C' of $\mathcal{G}^{\mathcal{A}}_{t,\{\rightsquigarrow\}}$ is inherited from the one of $\mathcal{A}$: $C'(x, q, \alpha) = C(q)$. A tree t is *accepted* by $\mathcal{A}$ with relation profile $\{\rightsquigarrow\}$ if Eve has a winning strategy in $\mathcal{G}^{\mathcal{A}}_{t,\{\rightsquigarrow\}}$, and we denote by $\mathcal{L}(\mathcal{A}, \{\rightsquigarrow\})$ the set of trees accepted by $\mathcal{A}$ equipped with relation profile $\{\rightsquigarrow\}$. If $\mathcal{A}$ is an alternating automaton (*i.e.* it only uses automata directions $\Diamond$ and $\Box$), it needs not be equipped by a relation profile to evaluate trees, and we write $\mathcal{L}(\mathcal{A})$ for the set of trees it accepts.

▶ Remark. In general, JTA can identify children of a given current node and send different copies independently to each one of them. This ability is not always needed, but quantifying (existentially or universally) over children is sufficient. This is the case in this work, reason why we have presented here a *symmetric* version of jumping tree automata, just like *symmetric* alternating automata have sometimes been considered (see *e.g.* [13]).

In the following, the size of a formula φ, written $|\varphi|$, is its number of subformulas, and the size of an automaton $\mathcal{A}$, written $|\mathcal{A}|$, is the size of its transition function (*i.e.* the sum of the sizes of formulas occuring in it).

3 Equivalence of jumping tree automata and epistemic μ-calculus

We show that JTA and L_μ^K are equally expressive, as stated by the following theorem.

▶ **Theorem 5.**

- *For every formula $\varphi \in \mathsf{L}_\mu^K$, there exists a jumping automaton $\mathcal{A}_\varphi$ such that for every relation profile $\{\rightsquigarrow\}$, $\mathcal{L}(\varphi, \{\rightsquigarrow\}) = \mathcal{L}(\mathcal{A}_\varphi, \{\rightsquigarrow\})$.*
- *For every jumping automaton $\mathcal{A}$, there exists an L_μ^K-formula $\varphi_\mathcal{A}$ such that for every relation profile $\{\rightsquigarrow\}$, $\mathcal{L}(\mathcal{A}, \{\rightsquigarrow\}) = \mathcal{L}(\varphi_\mathcal{A}, \{\rightsquigarrow\})$.*

Moreover, the translations are effective and linear.

The rest of this section is dedicated to the proof of Theorem 5 and to two corollaries.

We rely on the classical equivalence between the multi-modal μ-calculus, written here $\mathbb{L}_\mu$, and alternating tree automata, when interpreted over transition systems: A (multi-modal, AP-labelled) *transition system* is a tuple $\mathcal{S} = (Q, \{R_i\}_{i\in I}, V)$, where Q is a set of *states*, I is a finite set of indices, for each $i \in I$, $R_i \subseteq Q \times Q$ is a binary relation, and $V : Q \to 2^{AP}$ is a *labelling function*. We do not detail the semantics of the μ-calculus and alternating automata over transition systems, which is very similar to the one for trees (see [7, Chap. 10]).

▶ **Proposition 6.** *[7, Chap. 9, Chap. 10]*

- *For every formula $\varphi \in \mathbb{L}_\mu$, there exists an alternating automaton $\mathcal{A}_\varphi$ that accepts precisely the transition systems verifying φ.*
- *For every alternating automaton $\mathcal{A}$, there exists an $\mathbb{L}_\mu$-formula $\varphi_\mathcal{A}$ whose models are exactly the transition systems accepted by $\mathcal{A}$.*

Moreover, the translations are effective and linear.

Now we make observation that AP-trees are connected, acyclic, rooted transition systems with one relation. Also, an AP-tree $t = (\tau, \ell)$ together with a relation profile $\{\rightsquigarrow_i\}_{i\in Ag}$ over $(2^{AP})^*$ can be seen as a transition system $\mathcal{S}_t^{\{\rightsquigarrow\}} = (\tau, \{R\} \cup \{R_i\}_{i\in Ag}, \ell)$, where xRy if y is a child of x, and xR_iy if $x \rightsquigarrow_i y$. For a relation profile $\{\rightsquigarrow\}$, we define $\mathcal{C}^{AP}_{\{\rightsquigarrow\}} := \{\mathcal{S}_t^{\{\rightsquigarrow\}} \mid t \text{ is an } AP\text{-tree}\}$, the class of all transition systems obtained by combining $\{\rightsquigarrow\}$ with AP-trees. Now, two additional simple observations are necessary to prove Theorem 5: (1) Given a relation profile $\{\rightsquigarrow\}$, an L_μ^K-formula on AP-trees can be seen as an $\mathbb{L}_\mu$-formula on $\mathcal{C}^{AP}_{\{\rightsquigarrow\}}$, and (2) A jumping automaton equipped with a relation profile $\{\rightsquigarrow\}$ and working on AP-trees can be seen as an alternating automaton working on $\mathcal{C}^{AP}_{\{\rightsquigarrow\}}$.

We now argue for Theorem 5: For the first point, take a formula $\varphi \in \mathsf{L}_\mu^K$ and a relation profile $\{\rightsquigarrow\}$. See it as an $\mathbb{L}_\mu$-formula over $\mathcal{C}^{AP}_{\{\rightsquigarrow\}}$. By Proposition 6, one can build in linear time an alternating automaton $\mathcal{A}_\varphi$ that has the same language as φ on transition systems, and therefore also when restricted to $\mathcal{C}^{AP}_{\{\rightsquigarrow\}}$. This $\mathcal{A}_\varphi$, when restricted to $\mathcal{C}^{AP}_{\{\rightsquigarrow\}}$, can be seen as a jumping automaton. Because $\mathcal{A}_\varphi$ only depends on φ and not on $\{\rightsquigarrow\}$, we obtain the desired result. The second point of Theorem 5 is just dealt by rolling back the above argumentation.

Theorem 5 has two important corollaries. First, let us recall some definitions and results concerning recognizable relations and jumping automata. Let Σ be a finite alphabet.

▶ **Definition 7.** A relation $\rightsquigarrow \subseteq \Sigma^* \times \Sigma^*$ is *recognizable* if there are two families of regular languages $\mathcal{U}_1, \ldots, \mathcal{U}_n \subseteq \Sigma^*$ and $\mathcal{U}'_1, \ldots, \mathcal{U}'_n \subseteq \Sigma^*$ such that $\rightsquigarrow = \bigcup_{i=1}^{n} \mathcal{U}_i \times \mathcal{U}'_i$.

For example, epistemic relations of agents whose memory can be represented by finite state machines are recognizable relations (see [14]).

Given a recognizable relation $\rightsquigarrow$, one easily shows that the language $\{w\#w' \mid w \rightsquigarrow w'\}$ where $\#$ is a fresh symbol can be accepted by a finite-state word automaton; we let *size* of $\rightsquigarrow$, written $|\rightsquigarrow|$, is then the number of states of a minimal word automaton that recognizes the language $\{w\#w' \mid w \rightsquigarrow w'\}$.

▶ **Theorem 8.** *[15, 14] For every jumping automaton $\mathcal{A}$ equipped with a relation profile $\{\rightsquigarrow\}$, if every relation $\rightsquigarrow_i$ in $\{\rightsquigarrow\}$ is recognizable, then there is a two-way tree automaton $\mathcal{A}_{\{\rightsquigarrow\}}$ that accepts the same language, and such that $|\mathcal{A}_{\{\rightsquigarrow\}}|$ is polynomial in $|\mathcal{A}| + \sum_{i \in Ag} |\rightsquigarrow_i|$.*

Restricting attention to trees of bounded arity, we obtain the following two corollaries:

▶ **Corollary 9.** *The satisfiability problem for epistemic μ-calculus with recognizable relations is* Exptime*-complete.*

Proof. The upper bound follows from Theorem 5 together with Theorem 8 and the fact that, for trees of bounded arity, the emptiness problem for two-way tree automata is Exptime-complete [21]. The hardness follows from EXPTIME-hardness of the satisfiability problem for standard μ-calculus. ◀

▶ **Corollary 10.** *Epistemic μ-calculus with recognizable relations is not more expressive than (its fragment) the μ-calculus.*

Proof. By Propositions 6, it suffices to show that for each epistemic μ-calculus formula φ interpreted with recognizable relations, there exists an alternating tree automaton that accepts the models of φ. Let $\varphi \in \mathsf{L}_\mu^K$, and let $\{\rightsquigarrow\}$ be a relation profile of recognizable relations. By Theorem 5, there exists a jumping automaton $\mathcal{A}_\varphi$ such that $\mathcal{L}(\mathcal{A}_\varphi, \{\rightsquigarrow\}) = \mathcal{L}(\varphi, \{\rightsquigarrow\})$. Then, by Theorem 8, there is a two-way tree automaton $\mathcal{A}_\varphi^{\{\rightsquigarrow\}}$ such that $\mathcal{L}(\mathcal{A}_\varphi, \{\rightsquigarrow\}) = \mathcal{L}(\mathcal{A}_\varphi^{\{\rightsquigarrow\}})$. Finally, by [21], there is a non-deterministic (hence alternating) tree automaton $\mathcal{B}_\varphi^{\{\rightsquigarrow\}}$ such that $\mathcal{L}(\mathcal{B}_\varphi^{\{\rightsquigarrow\}}) = \mathcal{L}(\mathcal{A}_\varphi^{\{\rightsquigarrow\}})$, which concludes. ◀

4 Inexpressivity

In this section we prove the non-expressibility of ATL with imperfect information within the epistemic μ-calculus. We exhibit a formula of $\mathrm{ATL_i}$ and a relation profile that has no equivalent in the epistemic μ-calculus evaluated with the same relation profile.

Let $AP = \{p\}$, $Ag = \{a\}$ and $\mathrm{Act}_a = \mathrm{Act} = \{a_0, a_1\}$. We have $AP_{\mathrm{act}} = \{p_{a_0}, p_{a_1}\}$. Assume that Agent a is synchronous blindfold, *i.e.* she observes nothing but the occurence of moves. Her indistinguishability relation on (AP, Act)-tree arenas is therefore $\rightsquigarrow \subseteq (2^{AP \cup AP_{\mathrm{act}}})^*$, defined by $w \rightsquigarrow w'$ if $|w| = |w'|$. Consider the formula $\langle\!\langle a \rangle\!\rangle \mathbf{F} p \in \mathrm{ATL_i}(Ag)$. We prove that there is no formula of the epistemic μ-calculus that is equivalent to φ with regards to the singleton relation profile $\{\rightsquigarrow\}$. More formally:

▶ **Theorem 11.** *For all $\varphi' \in \mathsf{L}_\mu^K(AP \cup AP_{act}, Ag)$, $\mathcal{L}(\varphi', \rightsquigarrow) \neq \mathcal{L}(\langle\!\langle a \rangle\!\rangle \mathbf{F} p, Act, \rightsquigarrow)$.*

The rest of this section is dedicated to the proof of Theorem 11.

Assume towards a contradiction that there is a formula $\varphi' \in \mathsf{L}_\mu^K(AP \cup AP_{\mathrm{act}}, Ag)$ such that $\mathcal{L}(\varphi', \rightsquigarrow) = \mathcal{L}(\langle\!\langle a \rangle\!\rangle \mathbf{F} p, \mathrm{Act}, \rightsquigarrow)$. By Theorem 5, there is a jumping automaton $\mathcal{A}$ such that $\mathcal{L}(\varphi', \rightsquigarrow) = \mathcal{L}(\mathcal{A}, \rightsquigarrow)$. Let $\mathcal{A} = (AP \cup AP_{\mathrm{act}}, Q, \delta, q_0, C)$, and let $N = |Q| + 1$.

We build 2^N tree-arenas in which the formula $\langle\!\langle a \rangle\!\rangle \mathbf{F} p$ holds. In each of them, the objective $\mathbf{F} p$ is attained with a different uniform strategy. We exhibit, for each tree, a winning strategy in the acceptance game of $\mathcal{A}$ on that tree, and then we employ the "pigeon hole" principle to

show that at least two of these strategies can be combined into a new strategy that accepts a new tree-arena, in which the only strategy for a to ensure $\mathbf{F}p$ is not uniform.

We describe the family of tree-arenas that we consider (see Figure 1). Concretely we only describe finite trees, infinite trees are obtained by adding loops on leafs and unfolding the obtained graphs. For each $i \in \{1, \dots, 2^N\}$, the tree $t_i = (\tau_i, \ell_i)$ is such that:

1. The root does not verify p: $\ell_i(\epsilon) = \emptyset$
2. In ϵ, Agent a can only play a_0. Through this action she can move to 2^N+2 different children. The first 2^N ones verify p, but not the last two ones. Formally, $\tau_i \cap \mathbb{N} = \{0, \dots, 2^N + 1\}$. For readability, we call x_{m+1} the node m for each $m \in \{0, \dots, 2^N + 1\}$ (see Figure 1). For $1 \leq k \leq 2^N$, $\ell_i(x_k) = \{p, p_{a_0}\}$, and for $k \in \{2^N + 1, 2^N + 2\}$, $\ell_i(x_k) = \{p_{a_0}\}$.
3. For $1 \leq k \leq 2^N + 2$, node x_k has exactly one child $y_k = x_k \cdot 0$ reachable through a_0, where p does not hold: for $1 \leq k \leq 2^N + 2$, $\ell_i(y_k) = \{p_{a_0}\}$.
4. For each $k \leq 2^N + 2$, the subtree $t_i{\downarrow_{x_k}}$ is a full binary tree of height N in which each non-leaf node $x \succeq x_k$ has a left child, accessed through a_0, and a right child, accessed through a_1. The valuations are as follows. First, for the actions: for $1 \leq k \leq 2^N + 2$ and $w \in \{0,1\}^{\leq N}$, $p_{a_c} \in \ell_i(y_k \cdot w)$, where c is the last letter of w. Now, for the proposition p. For each $k \in \{1, \dots, 2^N\}$, let $w_k \in \{0,1\}^N$ be the binary representation of $k-1$. For $w \in \{0,1\}^{\leq N}$, if $1 \leq k \leq 2^N$, then $p \in \ell_i(y_k \cdot w)$ if and only if $w = w_k$, and if $k \in \{2^N + 1, 2^N + 2\}$, $p \in \ell_i(y_k \cdot w)$ if and only if $w = w_i$.

Observe that for all $i, j \in \{1, \dots, 2^N\}$, t_i and t_j share the same underlying tree, that we shall write τ: $\tau_i = \tau_j = \tau$. Moreover, the labellings only differ on the leafs of $\tau{\downarrow_{y_{2^N+1}}}$ and $\tau{\downarrow_{y_{2^N+2}}}$. Remark also that, since Agent a observes no atomic proposition, her uniform strategies are simply (infinite) sequences of actions. Also, for each i such that $1 \leq i \leq 2^N$, $\mathcal{G}^i$ denotes $\mathcal{G}^{\mathcal{A}}_{t_i,\{\rightsquigarrow\}}$, the acceptance game of $\mathcal{A}$ on t_i with relation $\rightsquigarrow$.

▶ **Lemma 12.** *For all $i \in \{1, \dots, 2^N\}$, Eve has a winning strategy in $\mathcal{G}^i$.*

Proof. Let $i \in \{1, \dots, 2^N\}$. Agent a has a uniform strategy in $\mathcal{G}^i$ for achieving $\mathbf{F}p$: it consists in playing $a_0 a_0 w_i a_0^\omega$. Therefore $t_i, \rightsquigarrow, \epsilon \models \langle\!\langle a \rangle\!\rangle \mathbf{F}p$, hence $t_i \in \mathcal{L}(\mathcal{A}, \rightsquigarrow)$. This precisely means that Eve has a winning strategy in $\mathcal{G}^i$. ◀

Let us take one winning strategy σ_i for Eve in each game $\mathcal{G}^i$. For each $1 \leq i \leq 2^N$, we define $\text{visit}_{\sigma_i} : \tau \to 2^Q$, which maps each node of τ to the set of states in which σ_i visits this node: $\text{visit}_{\sigma_i}(x) := \{q \mid \exists \pi \in \text{Out}(\sigma_i), \exists i \geq 0, \exists b \in \mathcal{B}^+(Dir \times Q) \text{ s.t. } \pi[i] = (x, q, b)\}$. Consider, for each $1 \leq i \leq 2^N$, the set $\text{visit}_{\sigma_i}(y_{2^N+1})$. Since there are at most $2^{|Q|}$ different such sets of states, and we have 2^N strategies with $N = |Q| + 1$, there must exist $i \neq j$ s.t. $\text{visit}_{\sigma_i}(y_{2^N+1}) = \text{visit}_{\sigma_j}(y_{2^N+1})$. For the rest of the proof we fix such a pair (i, j). We now consider the tree-arena t_0 that consists in t_i where the subtree $t_i{\downarrow_{y_{2^N+1}}}$ is replaced with $t_j{\downarrow_{y_{2^N+1}}}$ (see Figure 1). Let us write $\mathcal{G}^0$ for $\mathcal{G}^{\mathcal{A}}_{t_0,\{\rightsquigarrow\}}$.

Observe that the three games $\mathcal{G}^i$, $\mathcal{G}^j$ and $\mathcal{G}^0$ share the same set of positions: $V^0 = V^i = V^j = \tau \times Q \times \mathbb{B}^+(Dir \times Q) = V$. Also, for all $1 \leq k \leq 2^N + 2$, $\ell_0(y_k) = \ell_i(y_k) = \ell_j(y_k)$ $(= \{p_{a_0}\})$, that we now write ℓ. Because positions of the form $(y_k, q, \delta(q, \ell))$ play an important role in the following, we let $v_k^q := (y_k, q, \delta(q, \ell))$.

We first establish the following crucial lemma, which allows us to transfer the existence of winning strategies in positions v_k^q from $\mathcal{G}^i$ and $\mathcal{G}^j$ to $\mathcal{G}^0$ (see Appendix A for the proof).

▶ **Lemma 13.**

1. *For all $q \in Q$, for $k \neq 2^N + 1$, $\mathcal{G}^0, v_k^q \leftrightarrows \mathcal{G}^i, v_k^q$, and*
2. *for all $q \in Q$, for $k \neq 2^N + 2$, $\mathcal{G}^0, v_k^q \leftrightarrows \mathcal{G}^j, v_k^q$.*

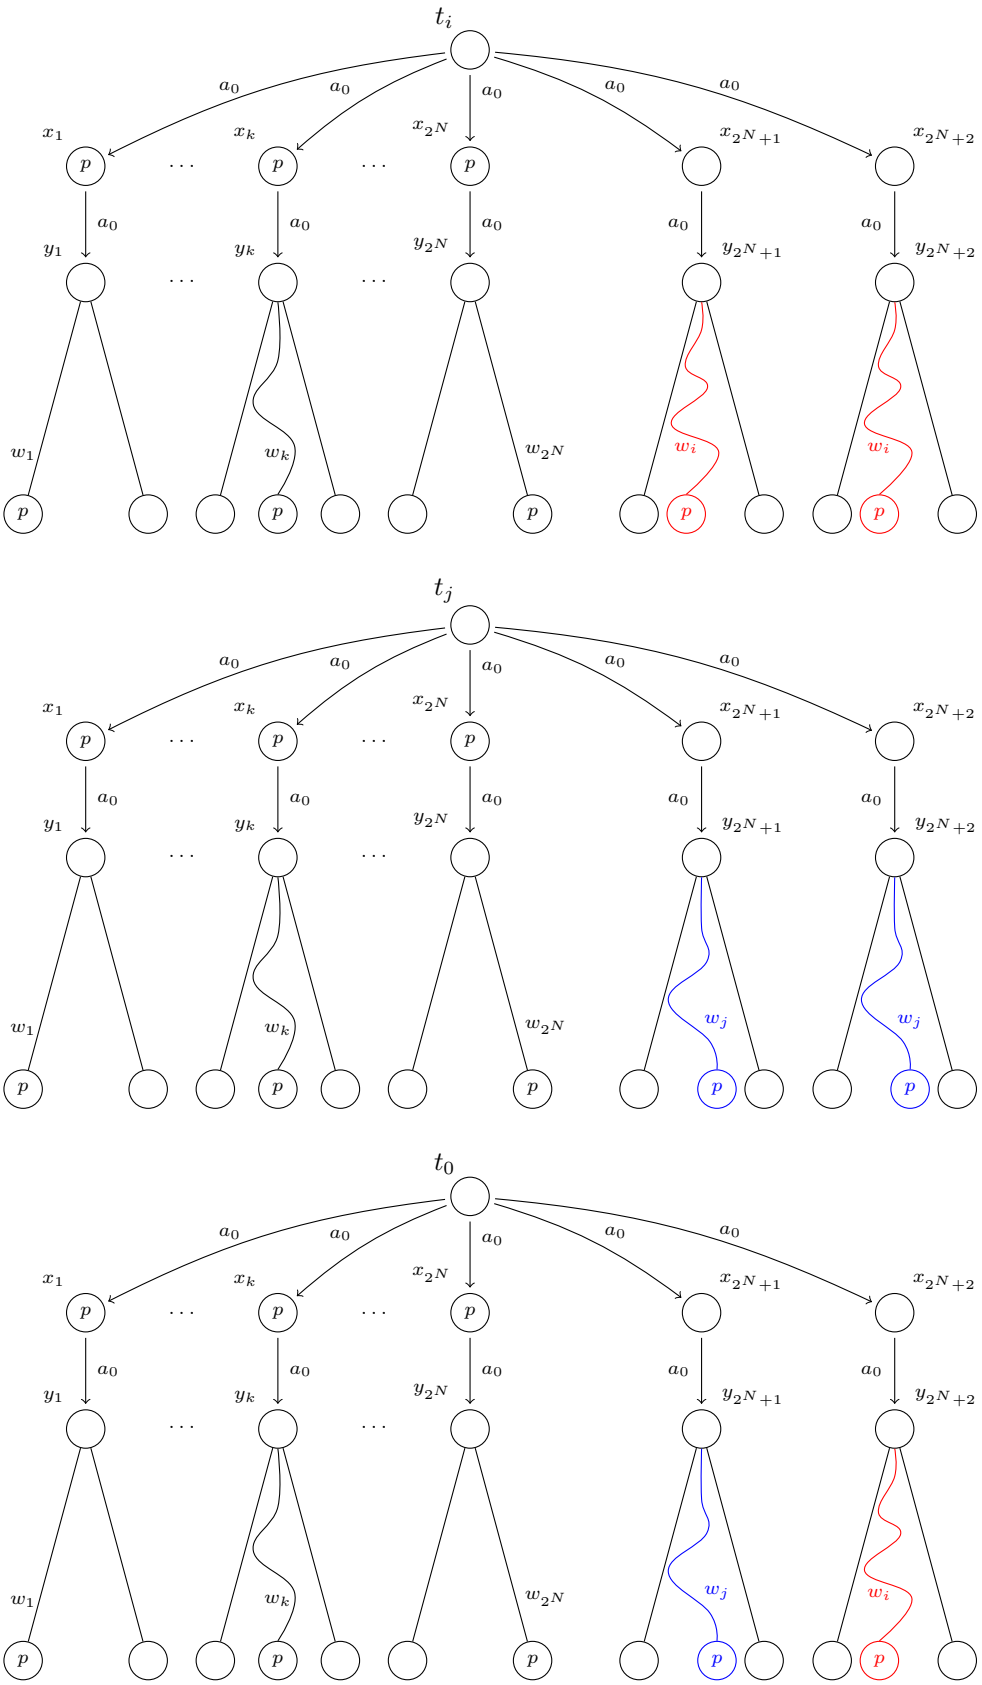

Figure 1 The tree t_i, the tree t_j, and the combined tree t_0.

Observe that, in t_0, Agent a has a non-uniform strategy to achieve $\mathbf{F}p$, but no uniform one. Therefore, $t_0, \rightsquigarrow, \epsilon \not\models \langle\langle a \rangle\rangle \mathbf{F}p$, and thence $t_0 \notin \mathcal{L}(\mathcal{A}, \rightsquigarrow)$. By definition of the acceptance for jumping automata, Eve does not have a winning strategy in $\mathcal{G}^0$. We prove the following proposition and obtain a contradiction, which terminates the proof of Theorem 11.

▶ **Proposition 14.** *Eve has a winning strategy in* $\mathcal{G}^0$.

Proof sketch. We give an intuition on how a winning strategy σ_0 for Eve in $\mathcal{G}^0$ can be obtained. The detail can be found in Appendix B. Let us define $\text{Start}_\tau = \{\epsilon, x_1, \ldots, x_{2^N+2}\}$, the two first levels of τ, and $\text{Start}_\mathcal{G} = \{(x, q, \alpha) \in V \mid x \in \text{Start}_\tau\}$. Observe that every play in $\mathcal{G}^0$ starts in $\text{Start}_\mathcal{G}$, namely, in $v_0 = (\epsilon, q_0, \delta(q_0, \ell_0(\epsilon)))$, and may remain in $\text{Start}_\mathcal{G}$ for an arbitrarily long time if it keeps jumping without going down. Otherwise, it exits $\text{Start}_\mathcal{G}$ by reaching some node y_k, in position v_k^q for some q. Observe also that from any position of $\text{Start}_\mathcal{G}$, the set of moves available in $\mathcal{G}^0$ and in $\mathcal{G}^i$ (and in $\mathcal{G}^j$) are the same. In $\mathcal{G}^0$, we let Eve follow σ_i as long as the game is in $\text{Start}_\mathcal{G}$. If the game remains in $\text{Start}_\mathcal{G}$ for ever, the obtained play is an outcome of σ_i which is winning for Eve in $\mathcal{G}^i$. Because a position has the same colour in all games, this play is also winning for Eve in $\mathcal{G}^0$. Otherwise, the game reaches a position of the form v_k^q. If $k \neq 2^N + 1$, because v_k^q has been reached by following σ_i which is winning in $\mathcal{G}^i$, v_k^q is a winning position for Eve in $\mathcal{G}^i$. By Point 1 of Lemma 13, $\mathcal{G}^0, v_k^q \leftrightarrows \mathcal{G}^i, v_k^q$, and by Proposition 2 we obtain that Eve also has a winning strategy from v_k^q in $\mathcal{G}^0$. If $k = 2^N + 1$, because $\text{visit}_{\sigma_i}(y_{2^N+1}) = \text{visit}_{\sigma_j}(y_{2^N+1})$, σ_j also visits position $v_{2^N+1}^q$, and therefore it is a winning position for Eve in $\mathcal{G}^j$. Again, by Point 2 of Lemma 13, $\mathcal{G}^0, v_k^q \leftrightarrows \mathcal{G}^j, v_k^q$, and by Proposition 2 Eve also has a winning strategy from v_k^q in $\mathcal{G}^0$. ◀

5 Conclusions

We have investigated in the expressive power of the epistemic μ-calculus by comparing it with jumping automata and ATL_i. For the first comparison, we have shown that, like in the classic case, L_μ^K is expressively equivalent to alternating jumping tree automata. Next, we have shown that ATL_i may express properties not expressible in L_μ^K, when interpreted with synchronous perfect-recall semantics. We have also shown that L_μ^K has a decidable satisfiability problem when the semantics relies on recognizable relations, *i.e.* bounded-memory semantics.

From the first two results above, one may prove that the monadic second order logic on trees, enriched with the equal-level predicate ($\text{MSO}_{eqlevel}$) [19], is strictly more expressive than L_μ^K: on the one hand, for each jumping automaton, one may build an equivalent $\text{MSO}_{eqlevel}$ formula, by appropriately encoding Eve's winning strategies in the automaton. On the other hand, it is not hard to see that $\text{MSO}_{eqlevel}$ may encode any ATL_i formula. These results strengthen the common belief that there exists no "fixpoint" axiomatization of ATL_i, contrary to what is known for ATL with perfect information, where the coalition operators have fixpoint expansions.

We plan to further investigate the impact of these results on a theory of jumping automata and their relation with MSO with the equal-level predicate, or other binary predicates. We conjecture that languages of jumping automata are not closed under existential quantifications. We also plan to identify a generalization of jumping automata which would be expressively equivalent (modulo bisimulations) to MSO with additional predicates. On the other hand, our non-expressiveness proof relies on the synchronous perfect recall setting, and we do not have an easy generalization to the case of non-synchronous perfect recall semantics, or to other types of semantics based on non-recognizable indistinguishability relations.

References

1 R. Alur, Th.A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.

2 D. Berwanger and L. Kaiser. Information tracking in games on graphs. *Journal of Logic, Language and Information*, 19(4):395–412, 2010.

3 N. Bulling, J. Dix, and W. Jamroga. Model checking logics of strategic ability: Complexity. In M. Dastani, K. V. Hindriks, and J.-J. C. Meyer, editors, *Specification and Verification of Multi-Agent Systems*, pages 125–160. Springer, 2010.

4 N. Bulling and W. Jamroga. Alternating epistemic mu-calculus. In *Proceedings of IJCAI'2011*, pages 109–114. IJCAI/AAAI, 2011.

5 Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.

6 E.A. Emerson. Handbook of theoretical computer science: Formal models and semantics, 1990.

7 E. Grädel, W. Thomas, and Th. Wilke. *Automata, Logics, and Infinite Games, volume 2500 of LNCS*. Springer Verlag, 2002.

8 J.Y. Halpern and M.Y. Vardi. The complexity of reasoning about knowledge and time. 1. Lower bounds. *Journal of Computer and System Sciences*, 38(1):195–237, 1989.

9 W. Jamroga and T. Ågotnes. What agents can achieve under incomplete information. In *Proceedings of AAMAS'2006*, pages 232–234. ACM, 2006.

10 W. Jamroga and Th. Ågotnes. Constructive knowledge: what agents can achieve under imperfect information. *Journal of Applied Non-Classical Logics*, 17(4):423–475, 2007.

11 D. Janin and I. Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *Proceedings of CONCUR'96*, pages 263–277. Springer, 1996.

12 D. Kozen. Results on the propositional mu-calculus. *Theor. Comput. Sci.*, 27:333–354, 1983.

13 O. Kupferman, M.Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *J. of the ACM*, 47(2):312–360, 2000.

14 B. Maubert. *Logical foundations of games with imperfect information: uniform strategies*. PhD thesis, Université de Rennes 1, 2014.

15 B. Maubert and S. Pinchinat. Jumping automata for uniform strategies. In *FSTTCS'13*, pages 287–298, 2013.

16 S. Pinchinat. A generic constructive solution for concurrent games with expressive constraints on strategies. In *Proceedings of ATVA'07*, pages 253–267, 2007.

17 P.-Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93, 2004.

18 N.V. Shilov and N.O. Garanina. Combining knowledge and fixpoints. Technical Report Preprint n.98, `http://www.iis.nsk.su/files/preprints/098.pdf`, A.P. Ershov Institute of Informatics Systems, Novosibirsk, 2002.

19 Wolfgang Thomas. Infinite trees and automaton-definable relations over omega-words. *Theor. Comput. Sci.*, 103(1):143–159, 1992.

20 W. van der Hoek and M. Wooldridge. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, 75(1):125–157, 2003.

21 M.Y. Vardi. Reasoning about the past with two-way automata. In *Proceedings of ICALP'98*, volume 1443 of *Lecture Notes in Computer Science*, pages 628–641, 1998.

A Proof of Lemma 13

▶ **Lemma 13.**
1. *For all $q \in Q$, for $k \neq 2^N + 1$, $\mathcal{G}^0, v_k^q \leftrightarrows \mathcal{G}^i, v_k^q$, and*
2. *for all $q \in Q$, for $k \neq 2^N + 2$, $\mathcal{G}^0, v_k^q \leftrightarrows \mathcal{G}^j, v_k^q$.*

Proof. For convenience, for $v, v' \in V$ and $k \in \{0, i, j\}$, we shall write $v \to^k v'$ if $(v, v') \in E^k$.

We start with point 1. Let us define the binary relation $Z \subseteq V^0 \times V^i$ as the smallest relation such that, for all $q \in Q$ and all $\alpha \in \mathbb{B}^+(Dir \times Q)$:

- $\forall k \neq 2^N + 1$, $\forall x \in \tau{\downarrow_{y_k}}$, $(x, q, \alpha) Z (x, q, \alpha)$,
- $\forall w \in \{0,1\}^*$, $(y_{2^N+1} \cdot w, q, \alpha) Z (y_j \cdot w, q, \alpha)$, and
- $\forall w \in \{0,1\}^*$, $(y_i \cdot w, q, \alpha) Z (y_{2^N+1} \cdot w, q, \alpha)$.

We prove that Z is a bisimulation between $\mathcal{G}^0$ and $\mathcal{G}^i$. Take $(v, v') \in Z$. By definition of Z, v and v' are on the horizontal line of y_k or below. Also, there are x, x', q and α such that $v = (x, q, \alpha)$ and $v' = (x', q, \alpha)$.

First, for colour harmony: by definition of the colours in acceptance games, it holds that $C(v) = C(q) = C(v')$.

Now, for Zig, take $u \in V$ such that $v \to^0 u$. According to the possible moves in the semantic games (see Section 2.4), this move is of one of the three following kinds:

1. it decomposes α without moving in the tree nor changing state,
2. it goes down to a child of x in a state q', or
3. it jumps to a node y such that $x \rightsquigarrow y$ in a state q'.

Case 1: We have $u = (x, q, \beta)$, where β is some subformula of α. According to the definition of semantic games, this move is also possible in $\mathcal{G}^i$: $v' \to^i u$. Therefore, we let $u' = u$. Because we have $(x, q, \alpha') Z (x', q, \alpha')$ for some $\alpha' = \alpha$, by definition of Z, it is true for all α', and in particular $(x, q, \beta) Z (x', q, \beta)$. Finally, $u Z u'$.

Case 2: We have $\alpha = \Diamond q'$ or $\alpha = \Box q'$, $u = (y, q', \delta(q', \ell_0(y)))$ for some child y of x; write $\beta := \delta(q', \ell_0(y))$ and $y := x \cdot c$, where $c \in \{0, 1\}$.

First, observe that by definition of Z, x and x' are at the same level ($|x| = |x'|$), and therefore if $x \cdot c$ exists in τ, so does $x' \cdot c$. It follows, by definition of semantic games, that $v' \to^i (x' \cdot c, q', \delta(q', \ell_i(x' \cdot c)))$ is a legal move in $\mathcal{G}^i$; write $\beta' := \delta(q', \ell_i(x' \cdot c))$ and $u' := (x' \cdot c, q', \beta')$.

We distinguish three possibilities again, according to the definition of Z and the fact that $(x, q, \alpha) Z (x', q, \alpha)$.

- $x = x'$. We have $y = x \cdot c = x' \cdot c$. By definition of Z, we obtain that $y \notin \tau{\downarrow_{y_{2^N+1}}}$, so that $\ell_0(y) = \ell_i(y)$. Therefore $\beta = \beta'$, and $u = u'$, which, by definition of Z, entails that $u Z u'$.
- $x = y_{2^N+1} \cdot w$ for some w. Because $v Z v'$, we have $x' = y_j \cdot w$. By observing t_0 and t_i, we obtain that $\ell_0(y_{2^N+1} \cdot w \cdot c) = \ell_i(y_j \cdot w \cdot c)$, so $\beta = \beta'$, and again, by definition of Z, $u Z u'$.
- $x = y_i \cdot w$ for some w. Because $v Z v'$, we have that $x' = y_{2^N+1} \cdot w$. Again, it holds that $\ell_0(y_i \cdot w \cdot c) = \ell_i(y_{2^N+1} \cdot w \cdot c)$, therefore $\beta = \beta'$, and by definition of Z, $u Z u'$.

Case 3: We have $\alpha = \diamondsuit q'$ or $\alpha = \boxminus q'$ for some q', $u = (y, q', \beta)$ for some $x \rightsquigarrow y$ and $\beta = \delta(q', \ell_0(y))$. By definition of Z, $|x| = |x'|$, and because Agent a is blind, the nodes reachable from x and x' through $\rightsquigarrow$ coincide (they are all the nodes at the same level). We therefore hace $|x| = |x'| = |y|$. We distinguish two cases.

- $y \in \tau{\downarrow_{y_k}}$ for some $k \neq 2^N + 1$: since $|x'| = |y|$, we have that $x' \rightsquigarrow y$, and therefore the move $v' \to^i (y, q', \delta(q', \ell_i(y))) = u'$ is legal in $\mathcal{G}^i$. Now, because $\ell_0(y) = \ell_i(y)$, $u = u'$, hence $u Z u'$.

- $y \in \tau{\downarrow_{y_{2^N+1}}}$: let $y = y_{2^N+1} \cdot w$ for some w. We have that $|y_j \cdot w| = |y_{2^N+1} \cdot w| = |y| = |x'|$, hence $x' \rightsquigarrow y_j \cdot w$, and therefore $v \rightarrow^i (y_j \cdot w, q', \delta(q', \ell_i(y_j \cdot w))) = u'$ is a valid move in $\mathcal{G}^i$. And because $\ell_0(y_{2^N+1} \cdot w) = \ell_i(y_j \cdot w)$, $\delta(q', \ell_i(y_j \cdot w)) = \beta$, and therefore uZu'.

For Zag, the proof is almost the same, making use of the third point in the the definition of Z instead of the second one for simulating the moves of $\mathcal{G}^i$ that jump in $\tau{\downarrow_{y_{2^N+1}}}$. So Z is a bisimulation between $\mathcal{G}^0$ and $\mathcal{G}^i$ and, clearly, for all $q \in Q$, for $k \neq 2^N + 1$, $(y_k, q, \delta(q, \ell(y_k)))Z(y_k, q, \delta(q, \ell(y_k)))$, *i.e.* $v_k^q Z v_k^q$, so that $\mathcal{G}^0, v_k^q \leftrightarrows \mathcal{G}^i, v_k^q$.

We turn to the proof of the second point in Lemma 13.

We define the following binary relation $Z' \subseteq V^0 \times V^j$, very similar to Z, as the smallest relation such that, for all $q \in Q$ and all $\alpha \in \mathbb{B}^+(Dir \times Q)$:

- $\forall k \neq 2^N + 1$, $\forall x \in \tau{\downarrow_{y_k}}$, $(x, q, \alpha)Z'(x, q, \alpha)$,
- $\forall w \in \{0,1\}^*$, $(y_{2^N+2} \cdot w, q, \alpha)Z'(y_i \cdot w, q, \alpha)$, and
- $\forall w \in \{0,1\}^*$, $(y_j \cdot w, q, \alpha)Z'(y_{2^N+2} \cdot w, q, \alpha)$.

The only difference is that now, the moves that must be avoided are those that jump in $\tau{\downarrow_{y_{2^N+2}}}$, which is the part that differs between t_0 and t_j. The rest of the proof is just the same as for the first point. ◀

B Proof of Proposition 14

▶ **Proposition 14.** *Eve has a winning strategy in $\mathcal{G}^0$.*

Proof. We define a strategy σ_0 for Eve in $\mathcal{G}^0$, and we prove that it is a winning strategy. First, for each position of the form v_k^q, if v_k^q is a winning position for Eve in $\mathcal{G}^0$, we pick a winning strategy for Eve in $(\mathcal{G}^0, v_k^q)$ that we call $\sigma_{v_k^q}$. Recall that $\text{Start}_\tau = \{\epsilon, x_1, \dots, x_{2^N+2}\}$ consists in the two first levels of τ, and $\text{Start}_\mathcal{G} = \{(x, q, \alpha) \in V \mid x \in \text{Start}_\tau\}$. Take a partial play ρ in $\mathcal{G}^0$ ending in a position of Eve.

- If $\rho \in \text{Start}_\mathcal{G}^*$, $\sigma_0(\rho) := \sigma_i(\rho)$.
- Otherwise, there exist ρ', k, q and ρ'' such that $\rho = \rho' \cdot v_k^q \cdot \rho''$ and $\rho' \in \text{Start}_\mathcal{G}^*$.
  - If v_k^q is a winning position for Eve in $\mathcal{G}^0$, $\sigma_{v_k^q}$ is defined, and we let $\sigma_0(\rho) := \sigma_{v_k^q}(v \cdot \rho'')$.
  - Otherwise, define $\sigma_0(\rho)$ arbitrarily.

▶ **Lemma 15.** *σ_0 is winning for Eve in $\mathcal{G}^0$.*

Let $\pi \in \text{Out}(\mathcal{G}^0, \sigma_0)$. If $\pi \in \text{Start}_\mathcal{G}^\omega$, then π is also a play in $\mathcal{G}^i$ that, moreover, follows σ_i, which is winning for Eve in $\mathcal{G}^i$, so π is winning for Eve in $\mathcal{G}^0$ (recall that positions have the same colours in the different acceptance games). Otherwise, there exist ρ, k, q and π' such that $\pi = \rho \cdot v_k^q \cdot \pi'$ and $\rho \in \text{Start}_\mathcal{G}^*$. Because $\rho \cdot v_k^q$ is a partial play in $\mathcal{G}^i$ that follows σ_i, which is winning for Eve in $\mathcal{G}^i$, v_k^q is a winning position in $\mathcal{G}^i$. We distinguish two cases.

- $k \neq 2^N + 1$: since v_k^q is a winning position for Eve in $\mathcal{G}^i$, by Lemma 13 and Proposition 2, v_k^q is also a winning position for Eve in $\mathcal{G}^0$.
- $k = 2^N + 1$: necessarily $q \in \text{visit}_{\sigma_i}(y_{2^N+1})$, and because $\text{visit}_{\sigma_i}(y_{2^N+1}) = \text{visit}_{\sigma_j}(y_{2^N+1})$, some outcome of σ_j in $\mathcal{G}^j$ visits v_k^q, which makes v_k^q a winning position for Eve in $\mathcal{G}^j$.

In both cases, $\sigma_{v_k^q}$ is defined, and by definition of σ_0, $v_k^q \cdot \pi' \in \text{Out}((\mathcal{G}^0, v_k^q), \sigma_{v_k^q})$. Because $\sigma_{v_k^q}$ is winning for Eve in $(\mathcal{G}^0, v_k^q)$, $v_k^q \cdot \pi'$ verifies the parity condition, and therefore also does $\pi = \rho \cdot v_k^q \cdot \pi'$. So π is winning for Eve, and we are done. ◀